

METHOD AND APPARATUS FOR SELECTIVELY ACCELERATING NETWORK COMMUNICATIONS

BACKGROUND OF THE INVENTION

5 1. Field of the Invention

The present invention relates generally to networking systems, and more specifically, to a method and system for selectively accelerating network communications.

10 2. Background of the Invention

Present-day network systems communicate through a variety of protocols and channels in order to interconnect computers, telephony devices and other systems that required data or voice communications. Quality of Service (QoS) is a designator that is used in network systems to assign or request desirable data transfer characteristics, such as delay and bandwidth characteristics for a given channel. Service quality can be assigned on a per-user basis to provide several levels of
20 interconnect performance conforming to desired performance levels. Users may be charged fees for different performance levels. For example, a business connection or Internet Service Provider (ISP) serving multiple users will have a higher desired

performance level than an individual residential customer, and the fees for such performance can be assigned accordingly.

QoS levels are typically set within a network by a configuration manager, which can be coupled to the network or coupled to a network component such as a router. The configuration manager is a program running on a computer that permits setting of network addresses such as Internet Protocol (IP) addresses, QoS requirements for a given connection between addresses and protocols to be used for communication between networked devices.

Within a wired network, QoS requirements related to TCP/IP networking performance are not as difficult to meet as with a satellite communications coupled network or with wireless networking systems. Communications using the transmission control protocol (TCP) require a one-to-one acknowledge response for each TCP segment (packet) communicated. The delays associated with satellite communications can make it impossible to efficiently utilize available bandwidth in the channel, since each end station must wait for responses for every transmitted TCP segment. Proprietary accelerated protocols can be used over the satellite portion of the network that alleviate the performance bottlenecks caused by acknowledgement delay, but use

of these protocols is very processor and memory intensive. In order to avoid higher costs of processing power and memory for temporarily holding TCP segments until their receipt is confirmed on the other side of a satellite communications channel and for performing packet processing, it is desirable to provide an alternative to a completely accelerated communications channel or an communications channel that uses no acceleration.

Therefore, it would be desirable to provide a system and method for selectively accelerating network communications, so that satellite channels and other communications channels with long delays can be used for network communications, while not requiring the processing power and memory required to accelerate transmissions for all communications through a channel.

SUMMARY OF THE INVENTION

The above objective of selectively accelerating network communications is achieved in a method of network communications and a network communications system. The method of network communications classifies a communications session, and selectively accelerates the communications session in conformity with the classification of the session. The network communications system includes a packet processor that classifies a communications session and selectively encapsulates and decomposes transmissions in an accelerated format in conformity with the session classification.

The foregoing and other objectives, features, and advantages of the invention will be apparent from the following, more particular, description of the preferred embodiment of the invention, as illustrated in the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram depicting a satellite network communications channel within which the present invention may be embodied.

Figure 2 is a pictorial diagram depicting a configuration manager configuration table in accordance with a preferred embodiment of the invention.

Figure 3 is a flowchart depicting operation of a packet processor in accordance with a preferred embodiment of the invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Referring now to the figures and in particular to **Figure 1**, a satellite network communications channel within which the present invention may be embodied is depicted in a block diagram. A first uplink/downlink includes a router **11**, a transmitter-receiver **12** and an antenna **13**. Router **11** is coupled to a network **10** which can interchange packets with the first uplink/downlink. A satellite **15** contains a transmitter-receiver, transponder or other suitable circuitry for receiving and transmitting information using an antenna **14**. The configuration of a separate transmitter-receiver **12** and router **11** is an illustrative configuration. Other configurations such as one that integrates a router with a transmitter-receiver are contemplated by the present invention and represent additional embodiments thereof.

Packets transmitted via first uplink/downlink to satellite **15** may be forwarded to a second uplink/downlink that includes a router **18**, a transmitter-receiver **17** and an antenna **16**. Router **18** is coupled to a network **19** which can interchange packets with the second uplink/downlink. While the illustrative embodiment described in this section uses TCP/IP protocol, it will be understood by those of ordinary skill in the art that the

present invention may be used with other protocols, and that the advantages of the present invention are particularly applicable to those protocols that require the uplink to maintain storage of transmitted packets until a confirmation that the packet has
5 been received by the downlink side of the communications channel. Additionally, while the illustrative embodiment is directed to communications channels involving a satellite uplink/downlink it will be understood by those of ordinary skill in the art that the present invention may be used with other communications channels , and that the advantages of the present invention are particularly applicable to those communications channels that have a long transmission/reception delay.

Typically, satellite links for networks using transmission control protocol (TCP) must wait for an acknowledgement to be received for each TCP segment transmitted through the link. For example, a TCP segment transmitted from network **10** to network **19** will be acknowledged at network **10** after a delay greater than twice the propagation delay from antenna **13** to antenna **16**. The
20 additional delay over propagation delay is produced by the satellite systems, as well as router delays in the first and second uplink/downlink. The resulting transmission delay can be obviated by using a proprietary protocol over the satellite communications channel that does not require a one-to-one packet

acknowledgment, yet provides guaranteed packet delivery.

However, processing power required within router **11** and router **18** and memory requirements become excessive. Memory requirements are increased, since each router must maintain tables of packets pending acknowledgement via the proprietary protocol from the other end of the satellite communications channel.

The present invention overcomes these limitations by selectively accelerating some sessions to provide higher quality of service (QoS). While acceleration may be turned off when memory or processing requirements are taxed, this produces a random performance per session. It is desirable to provide connections with consistently high QoS for some users, who will pay for the service accordingly. It is also desirable to provide connections with a lower QoS for users who are more cost-conscious. The lower QoS connections may be upgraded when the communications channel is not being taxed by high demands from high performance connections, or the performance may be permanently fixed at a lower level.

Referring now to **Figure 2**, a configuration table that is managed by a configuration manager in accordance with a preferred embodiment of the invention is shown. The

configuration manager is a program that may configure the operation of router **11** or router **18** via standard network management protocol (SNMP), a console interface via a direct serial or Ethernet connection, via a TELNET session, or via a web-based interface accessed from a management IP address. The packet processor makes decisions in accordance with configuration tables produced by the configuration manager. These tables allow the characterization of sessions by source IP address, source port number, destination IP address and destination port number. The IP addresses may be a fixed number, range of addresses or a wildcard indicating all possible values, wherein the masks for ranges are indicated by a "/" and wildcards are indicated by an "*". Port numbers may be fixed numbers or wildcards. Transmission types are determined by the destination port number in accordance with RFC 1700, a specification maintained by the Network Working Group. For example, a destination port number of 20 indicates an FTP session, while a destination port number of 80 indicates an HTTP session.

The left-hand side of the configuration table depicts a table entry number, with higher priority assignments at the top of the table (lower entry numbers). Exemplary table entry 1 configures the packet processor such that up to five ftp

sessions originating from source IP address range 172.16.0.0/16 will be accelerated with a maximum rate per session of 512 kilobits per second for each transmission. Once five sessions have been established from this IP address range, any new sessions that are established from this IP address range will be not be accelerated until one of the five existing sessions has been terminated, thereby limiting the allocation of acceleration processing power and memory for this IP address.

The data rates indicated in the configuration table are approximate maximum levels of performance and are tailored by setting the buffer size for the session. Since buffering is required to overcome the delays associated with propagating packets through a satellite communications channel, controlling buffer size for a session effectively controls the throughput. The final column entry for controlling sessions originating from the specified source IP address range is whether or not resources are pre-allocated for the specified address range. If the resources are pre-allocated, sessions originating from the specified address range are guaranteed their QoS since the resources will always be available. The pre-allocation can be implemented by permanently reserving resources within the system based on the configuration table. Alternatively a higher priority session does not have associated pre-allocated

resources may be accelerated by pre-empting resources from a lower priority session that is already in progress.

Exemplary table entry 2 in **Figure 2** shows a configuration in which the packet processor will attempt to accelerate up to 10 HTTP sessions terminating in destination address range 10.0.0.0/24 are accelerated to a maximum data rate of 256 kilobits per second per session. Exemplary table entry 3 will attempt to accelerate all sessions originating from IP address range 192.168.10.0/8 subject to resource availability.

Exemplary table entry 4 in **Figure 2** is a "catch-all" configuration for FTP sessions. The system will attempt to accelerate all FTP sessions subject to resource availability and exemplary table entry 5 is a catch-all for all TCP sessions, causing all TCP sessions to be accelerated subject to remaining resources.

Referring now to **Figure 3**, a flowchart of a method of packet processing in a network is depicted in accordance with a preferred embodiment of the invention. The packet processor receives an IP datagram (**step 30**). The packet processor then checks the IP protocol field to determine if TCP is the higher level protocol (per RFC 1700, the TCP protocol number is 6). If

the protocol is not TCP (**decision 31**), the datagram is dispatched for normal processing (**step 38**). Next, the packet processor examines the source IP address, destination IP address, the source TCP port number and the destination TCP port number contained within the IP datagram. If a configuration table match is found (**decision 32**), the IP datagram is selected for further processing, otherwise the IP datagram is transmitted without acceleration (e.g., normal bypass processing) (**step 38**). If the number of accelerated sessions for that classification has reached a maximum threshold value (**decision 33**), the packet is likewise processed without acceleration (**step 38**). Finally, if there are no available resources (**decision 34**) and there is not an existing lower priority session (**decision 36**), the packet is also processed without acceleration (**step 38**). Otherwise, if resources are available (**decision 34**) the packet is processed using an acceleration mechanism (**step 35**) that terminates the TCP session locally, converts the data to a proprietary protocol for transmission over the satellite communications channel and recreates the TCP session at the other side of the satellite communication for transmission to the attached network, to provide accelerated packet transmission through the channel. If resources are not available (**decision 34**) but a lower priority session exists (**decision 36**), then resources are pre-empted from the lower priority session (**step 37**) and the packet is likewise

accelerated (**step 35**). The above method is an illustrative embodiment and other variations may be performed which conform to preferred embodiments of the present invention. For example, the decision tree for packet processing may be performed in
5 different orders without departing from the scope of the present invention.

While the invention has been particularly shown and described with reference to the preferred embodiments thereof,
10 it will be understood by those skilled in the art that the foregoing and other changes in form, and details may be made therein without departing from the spirit and scope of the invention.